arXiv:1608.07505v1 [cs.DS] 26 Aug 2016

A Note on the Practicality of Maximal Planar Subgraph Algorithms

Markus Chimani[1], Karsten Klein[2], Tilo Wiedera[1]

[1] University Osnabrück, Germany
{markus.chimani,tilo.wiedera}@uni-osnabrueck.de
[2] Uni Konstanz, Germany
karsten.klein@uni-konstanz.de

Abstract. Given a graph G, the NP-hard *Maximum Planar Subgraph* problem (*MPS*) asks for a planar subgraph of G with the maximum number of edges. There are several heuristic, approximative, and exact algorithms to tackle the problem, but—to the best of our knowledge—they have never been compared competitively in practice.
We report on an exploratory study on the relative merits of the diverse approaches, focusing on practical runtime, solution quality, and implementation complexity. Surprisingly, a seemingly only theoretically strong approximation forms the building block of the strongest choice.

1 Introduction

We consider the problem of finding a large planar subgraph in a given non-planar graph $G = (V, E)$; $n := |V|$, $m := |E|$. We distinguish between algorithms that find a *large*, *maximal*, or *maximum* such graph: while the latter (MPS) is one with largest edge cardinality and NP-hard to find [18], a subgraph is inclusionwise maximal if it cannot be enlarged by adding any further edge of G. Sometimes, the inverse question—the *skewness* of G—is asked: find the smallest number $skew(G)$ of edges to remove, such that the remaining graph is planar.

The problem is a natural non-trivial graph property, and the probably best known non-planarity measure next to crossing number. This already may be reason enough to investigate its computation. Moreover, MPS/skewness arises at the core of several other applications: E.g., the practically strongest heuristic to draw G with few crossings—the *planarization method* [2, 7][3]—starts with a large planar subgraph, and then adds the other edges into it.

Recognizing graphs of small skewness also plays a crucial role in parameterized complexity: Many problems become easier when considering a planar graph; e.g., maximum flow can be computed in $\mathcal{O}(n \log n)$ time, the Steiner tree problem allows a PTAS, the maximum cut can be found in polynomial time, etc. It hence can be a good idea to (in a nutshell) remove a couple of edges to obtain a planar graph, solve the problem on this subgraph, and then consider suitable

[3] In contrast to this meaning, the MPS problem itself is also sometimes called *planarization*. We refrain from the latter use to avoid confusion.

modifications to the solution to accommodate for the previously ignored edges. E.g., we can compute a maximum flow in time $\mathcal{O}(skew(G)^3 \cdot n \log n)$ [13].

While solving MPS is NP-hard, there are diverse polynomial-time approaches to compute a large or maximal planar subgraph, ranging from very trivial to sophisticated. By Euler's formula we know that already a spanning tree gives a 1/3-approximation for MPS. Hence all reasonable algorithms achieve this ratio. Only the *cactus algorithms* (see below) are known to exhibit better ratios. We will also consider an exact MPS algorithm based on integer linear programs (ILPs).

All algorithms considered in this paper are known (for quite some time, in fact), and are theory-wise well understood both in terms of worst case solution quality and running time. To our knowledge, however, they have never been practically compared. In this paper we are in particular interested in the following quality measures, and their interplay:

– What is the practical difference in terms of running time?
– What is the practical difference in solution quality (i.e., subgraph density)?
– What is the implementation effort of the various approaches?

Overall, understanding these different quality measures as a multi-criteria setting, we can argue for each of the considered algorithms that it is pareto-optimal. We are in particular interested in studying a somewhat "blurred" notion of pareto-optimality: We want to investigate, e.g., in which situations the additional sophistication of algorithms gives "significant enough" improvements.[4]

Also the measure of "implementation complexity" is surprisingly hard to concisely define, and even in the field of software-engineers there is no prevailing notion; "lines of code" are, for example, very unrelated to the intricacies of algorithm implementation. We will hence only argue in terms of direct comparisons between pairs of algorithms, based on our experience when implementing them.[5]

As we will see in the next section, there are certain building blocks all algorithms require, e.g., a graph data structure and (except for `C`, see below) an efficient planarity test. When discussing implementation complexity, it seems safe to assume that a programmer will already start off with some kind of graph library for her basic datastructure needs.[6] In the context of the ILP-based approach, we assume that the programmer uses one of the various freely available (or commercial) frameworks. Writing a competitive branch-and-cut framework from ground up would require a staggering amount of knowledge, experience, time, and finesse. The ILP method is simply not an option if the programmer may not use a preexisting framework.

In the following section, we discuss our considered algorithms and their implementation complexity. In Section 3, we present our experimental study. We

[4] Clearly, there is a systematic weakness, as it may be highly application-dependent what one considers "significant enough". We hence cannot universally answer this question, but aim to give a guideline with which one can come to her own conclusions.

[5] This measure is still intrinsically subjective (although we feel that the situation is quite clear in most cases), and opinions may vary depending on the implementor's knowledge, experience, etc. We discuss these issues when they become relevant.

[6] Many freely available graph libraries contain linear-time planarity tests. They usually lack sophisticated algorithms for finding large planar subgraphs.

first consider the pure problem of obtaining a planar subgraph. Thereafter, we investigate the algorithm choices when solving MPS as a subproblem in a typical graph drawing setting—the planarization heuristic.

2 Algorithms

Naïve Approach (`Nï`). The algorithmically simplest way to find a maximal planar subgraph is to start with the empty graph and to insert each edge (in random order) unless the planarity test fails. Given an $\mathcal{O}(n)$ time planarity test (we use the algorithm by Boyer and Myrvold [3], which is also supposed to be among the practically fastest), this approach requires $\mathcal{O}(nm)$ overall time.[7]

In our study, we consider a trivial multi-start variant that picks the best solution after several runs of the algorithm, each with a different randomized order. The obvious benefit of this approach is the fact that it is trivial to understand and implement—once one has any planarity test as a black box.

Augmented Planarity Test (`BM`, `BM+`). Planarity tests can be modified to allow the construction of large planar subgraphs. We will briefly sketch these modifications in the context of the above mentioned $\mathcal{O}(n)$ planarity test by Boyer and Myrvold [3]: In the original test, we start with a DFS tree and build the original graph bottom-up; we incrementally add a vertex together with its DFS edges to its children and the backedges from its decendents. The test fails if at least one backedge cannot be embedded.

We can obtain a large (though in general not maximal) planar subgraph by ignoring whether some backedges have not been embedded, and continuing with the algorithm (`BM`). If we require maximality, we can use `Nï` as a prostprocessing to grow the obtained graph further (`BM+`). While this voids the linear runtime, it will be faster than the pure naïve approach. Given an implementation of the planarity testing algorithm, the required modifications are relatively simple per se—however, they are potentially hard to get right as the implementor needs to understand side effects within the complex planarity testing implementation.

Alternatively, Hsu [14] showed how to overcome the lack of maximality directly within the planarity testing algorithm [19] (which is essentially equivalent to [3]), retaining linear runtime. While this approach is the most promising in terms of running time, it would require the most demanding implementation of all approaches discussed in this paper (including the next subsection)—it means to implement a full planarity testing algorithm plus intricate additional procedures. We know of no implementation of this algorithm.

Cactus Algorithm (`C`, `C+`). The only non-trivial approximation ratios are achieved by two cactus-based algorithms [4]. Thereby, one starts with the disconnected vertices of G. To obtain a ratio of $7/18$ (`C`), we iteratively add triangles connecting formerly disconnected components. This process leaves a forest F of tree-like

[7] We *could* speed-up the process in practice by starting with a spanning tree plus three edges. However, there are instances where the initial inclusion of a spanning tree prohibits the optimal solution and restricts one to approximation ratios $\leq 2/3$ [8].

structures made out of triangles—*cactusses*. Finally, we make F connected by adding arbitrary edges of E between disconnected components. Since this subgraph will not be maximal in general, we can use `Nï` to grow it further (`C+`).

From the implementation point of view, this algorithm is very trivial and—unless one requires maximality—does not even involve any planarity test. While a bit more complex than the naïve approach, it does not require modifications to complex and potentially hard-to-understand planarity testing code like `BM`.

For the best approximation ratio of 4/9 one seeks not a maximal but a maximum cactus forest. However, this approach is of mostly theoretical interest as it requires non-trivial polynomial time matroid algorithms.

ILP Approach (`ILP`). Finally, we use an integer linear program (ILP) to solve MPS exactly in reasonable (but formally non-polynomial) time, see [15]. With binary variables for each edge, specifying whether it is in the solution, we have

$$\max \left\{ \sum\nolimits_{e\in E} x_e \;\middle|\; \sum\nolimits_{e\in K} x_e \leq |K|-1 \text{ for all Kuratowski subdivisions } K \subseteq G \right\}.$$

Kuratowski's theorem [17] states that a graph is planar if and only if it does not contain a K_5 or a $K_{3,3}$ as a subdivision—*Kuratowski* subdivisions. Hence we guarantee a planar solution by requiring to remove at least one edge in each such subgraph K. While the set of these constraints is exponential in size, we can separate them heuristically within a branch-and-cut framework, see [15]: after each LP relaxation, we round the fractional solution and try to identify a Kuratowski subdivision that leads to a violated constraint.

This separation in fact constitutes the central implementation effort. Typical planarity testing algorithms initially only answer *yes* or *no*. In the latter case, however, all known linear-time algorithms can be extended to extract a *witness* of non-planarity in the form of a Kuratowski subdivision in $\mathcal{O}(n)$ time. If the available implementation does not support this additional query, it can be simulated using $\mathcal{O}(n)$ calls to the planarity testing algorithm, by incrementally removing edges whenever the graph stays non-planar after the removal. Both methods result in a straight-forward implementation (assuming some familiarity with ILP frameworks), but an additional tuning step to decide, e.g., on rounding thresholds, is necessary. The overall complexity is probably somewhere in-between `C` and `BM`. In our study, we decided to use the effective extraction scheme described in [10] which gives several Kuratowski subdivions via a single call. We propose, however, to use this feature only if it is already available in the library: its implementation effort would otherwise be comparable to a full planarity test, and in particular for harder instances its benefit is not significant.

3 Experiments

For an exploratory study we conducted experiments on several benchmark sets. We summarize the results as follows—observe the inversion between *F1* and *F2*.

F1. `C+` yields the best solutions. Choosing a "well-growable" initial subgraph—in our case a good cactus—is practically important. The better solution of `BM` is a weak starting point for `BM+`; even `Nï` gives clearly better solutions.

F2. `BM` gives better solutions than `C`; both are the by far fastest approaches. Thus, if runtime is more crucial than maximality, we suggest to use `BM`.
F3. `ILP` only works for small graphs. Expander graphs (they are sparse but well-connected) seem to be among the hardest instances for the approach.
F4. Larger planar subgraphs lead to fewer crossings for the planarization method. However, this is much less pronounced with modern insertion methods.

Setup & Instances. All considered algorithms are implemented in C++ (g++ 5.3.1, 64bit, `-O3`) as part of OGDF [5], the ILP is solved via CPLEX 12.6. We use an Intel Xeon E5-2430 v2, 2.50 GHz running Debian 8; algorithms use singles cores out of twelve, with a memory limit of 4GB per process.

We use the non-planar graphs of the established benchmark sets North [12] (423 instances), Rome [11] (8249), and SteinLib [16] (586), all of which include real-world instances. In our plots, we group instances according to $|V|$ rounded to the nearest multiple of 10; for Rome we only consider graph with ≥ 25 vertices.

Additionally, we consider two artificial sets: BaAl [1] are scale-free graphs, and Regular [20] (implemented as part of the OGDF) are random regular graphs; they are *expander graphs* w.h.p. [folklore]. Both sets contain 20 instances for each combination of $|V| \in \{10^2, 10^3, 10^4\}$ and $|E|/|V| \in \{2, 3, 5, 10, 20\}$.

Evaluation. Our results confirm the need for heuristic approaches, as `ILP` solves less than 25% of the larger graphs of the (comparably simple) Rome set within 10min. Even deploying strong preprocessing [6] (`+PP`) and doubling the computation time does not help significantly, cf. Fig. 1(d). Already 30-vertex graphs with density 3, generated like Regular, cannot be solved within 48 hours (→*F3*).

We measure solution quality by the *density* (edges per vertices) of the computed planar subgraph. Independently of the benchmark set, `C+` always achieves the best solutions, cf. Fig. 1(a)(b)(table) (→*F1*). We know instances where `Nï` is only a 1/3 approximation whereas the worst ratio known for `BM+` is 2/3 [8]. Surprisingly, `Nï` yields distinctly better solutions than `BM+` in practice (→*F1*).

On SteinLib, BaAl, and Regular, both `C` and `BM` behave similar w.r.t. solution quality. For Rome and North, however, `BM` yields solutions that are 20–30% better, respectively (→*F2*). This discrepancy seems to be due the fact that the found subgraphs are generally very sparse for both algorithms on BaAl and Regular (average density of 1.1 and 1.2, respectively, for the largest graphs).

Both `C` and `BM` are extremly (and similarly) fast; Fig. 1(c)(table) (→*F2*). For `BM+` and `C+`, the `Nï`-postprocessing dominates the running time: `Nï` is worst, followed by `C+` and `BM+`—a larger initial solution leads to fewer trys for growing. Nonetheless, we observe that the (weaker) solution of `C` allows for significantly more *successful* growing steps that `BM` (→*F1*).

Finally, we investigate the importance of the subgraph selection for the planarization method, cf. Fig. 1(e)(f). For the simplest insertion algorithms (iterative edge insertions, fixed embedding, no postprocessing, [2]), a strong subgraph method (`C+`) is important; `C` leads to very bad solutions. For state-of-the-art insertion routines (simultaneous edge insertions, variable embedding, strong postprocessing, [7, 9]) the subgraph selection is less important; even `C` is feasible.

	density, relative to best										runtime [s]									
	StеinLib				BaAl			Regular			StеinLib				BaAl			Regular		
	B-E	I*	S†	V†	2	3	4	2	3	4	B-E	I*	S	V	2	3	4	2	3	4
BM	.86	.85	.82	.84	.72	.64	.66	.82	.84	.95	.06	.07	.01	.11	.00	.02	1	.00	.02	2
BM+	.90	.90	.88	.86	.85	.74	.73	.89	.89	.97	47.34	8.85	2.70	90.97	.06	10.57	143	.03	4.60	236
C	.84	.67	.81	.87	.60	.68	.75	.74	.88	.95	.04	.04	.01	.07	.00	.01	0	.00	.01	1
C+	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	9.19	16.22	3.13	16.35	.06	12.32	152	.04	5.34	217
Nï	.92	.98	.91	.91	.96	.94	.92	.92	.91	.97	49.38	28.59	2.43	95.08	.05	7.92	239	.04	4.85	252

†S (= constr. sparse): PUC, SP; V (= VLSI): ALUE, ALUT, LIN, TAQ, DIW, DMXA, GAP, MSM, 1R, 2R

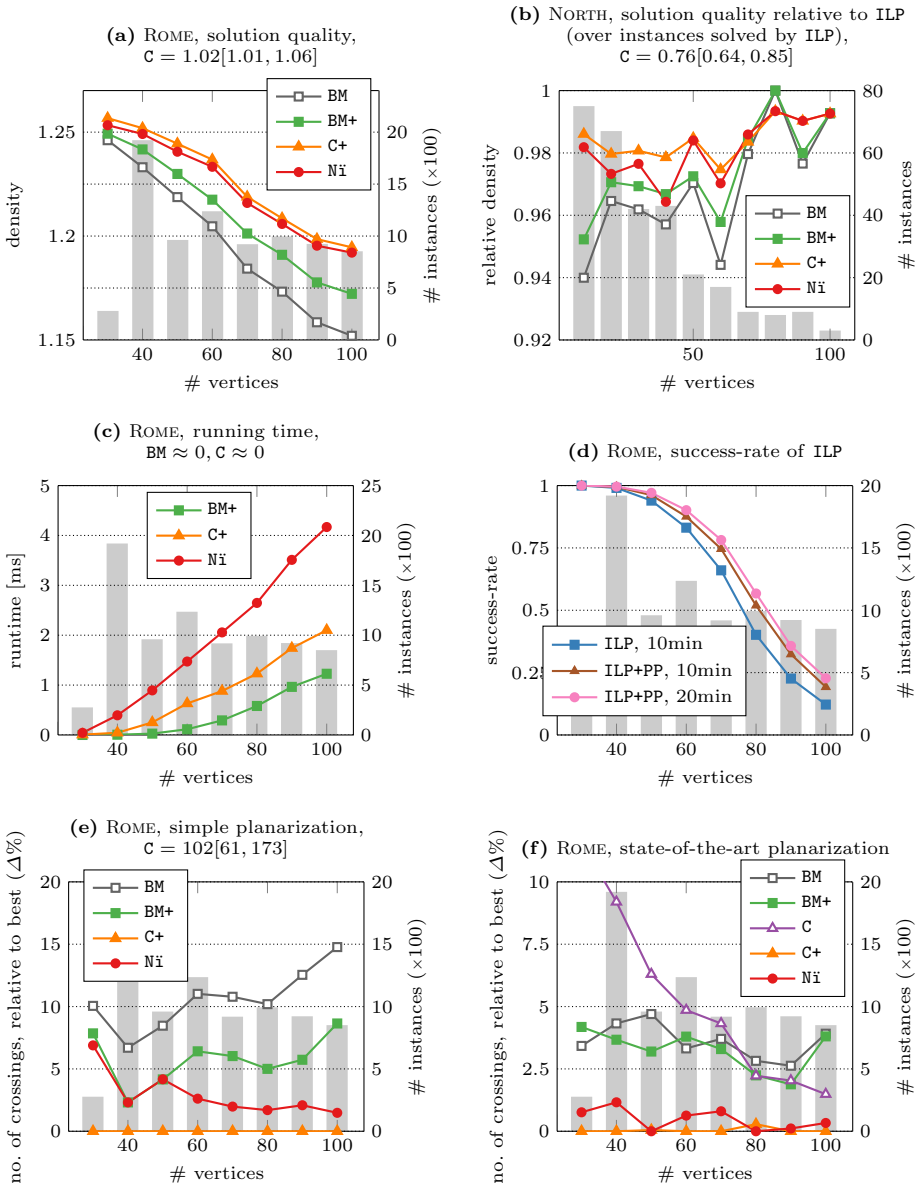

Fig. 1. We may omit algorithms whose values are unsuitable for a plot; instead we give their average[min, max] in the caption.

References

1. `BarabasiAlbertGenerator` of the java universal network/graph framework, `jung.sourceforge.net`
2. Batini, C., Talamo, M., Tamassia, R.: Computer aided layout of entity relationship diagrams. Journal of Systems and Software 4(2-3), 163–173 (1984)
3. Boyer, J.M., Myrvold, W.J.: On the cutting edge: Simplified $\mathcal{O}(n)$ planarity by edge addition. J. Graph Algorithms Appl. 8(2), 241–273 (2004)
4. Călinescu, G., Fernandes, C., Finkler, U., Karloff, H.: A better approximation algorithm for finding planar subgraphs. J. Algorithms 27, 269–302 (1998)
5. Chimani, M., Gutwenger, C., Jünger, M., Klau, G.W., Klein, K., Mutzel, P.: The open graph drawing framework (OGDF). In: Tamassia, R. (ed.) Handbook of Graph Drawing and Visualization, chap. 17, pp. 543–569. CRC Press (2014), `www.ogdf.net`
6. Chimani, M., Gutwenger, C.: Non-planar core reduction of graphs. Discrete Mathematics 309(7), 1838–1855 (2009)
7. Chimani, M., Gutwenger, C.: Advances in the planarization method: Effective multiple edge insertions. J. Graph Algorithms Appl. 16(3), 729–757 (2012)
8. Chimani, M., Hedtke, I., Wiedera, T.: Limits of greedy approximation algorithms for the maximum planar subgraph problem. In: Proc. Int'l Workshop on Combinatorial Algorithms (IWOCA) 2016. LNCS (2016), to appear
9. Chimani, M., Hlinený, P.: Inserting multiple edges into a planar graph. CoRR abs/1509.07952 (2015), `http://arxiv.org/abs/1509.07952`
10. Chimani, M., Mutzel, P., Schmidt, J.M.: Efficient extraction of multiple kuratowski subdivisions. In: Graph Drawing, 15th Int'l Symposium, GD 2007, Sydney, Australia, September 24-26, 2007. Revised Papers. pp. 159–170 (2007)
11. Di Battista, G., Garg, A., Liotta, G., Tamassia, R., Tassinari, E., Vargiu, F.: An experimental comparison of four graph drawing algorithms. Computational Geometry 7(56), 303–325 (1997)
12. Di Battista, G., Garg, A., Liotta, G., Parise, A., Tassinari, R., Tassinari, E., Vargiu, F., Vismara, L.: Drawing directed acyclic graphs: An experimental study. Int'l J. of Computational Geometry & Applications 10(06), 623–648 (2000)
13. Hochstein, J.M., Weihe, K.: Maximum s-t-flow with k crossings in $\mathcal{O}(k^3 n \log n)$. In: Proc. SODA 2007. pp. 843–847. ACM-SIAM (2007)
14. Hsu, W.: A linear time algorithm for finding a maximal planar subgraph based on PC-trees. LNCS, vol. 3595, pp. 787–797. Springer (2005)
15. Jünger, M., Mutzel, P.: Maximum planar subgraphs and nice embeddings: Practical layout tools. Algorithmica 16(1), 33–59 (1996)
16. Koch, T., Martin, A., Voß, S.: SteinLib: An updated library on steiner tree problems in graphs. Tech. Rep. ZIB-Report 00-37, Konrad-Zuse-Zentrum für Informationstechnik Berlin (2000), `http://elib.zib.de/steinlib`
17. Kuratowski, C.: Sur le problme des courbes gauches en topologie. Fundamenta Mathematicae 15(1), 271–283 (1930)
18. Liu, P.C., Geldmacher, R.C.: On the deletion of nonplanar edges of a graph. In: Proc. 10th Southeastern Conf. on Combinat., Graph Theory and Comput. pp. 727–738. Congress. Numer., XXIII–XXIV, Utilitas Math., Winnipeg, Man. (1979)
19. Shih, W., Hsu, W.: A new planarity test. Theor. Comput. Sci. 223(1-2), 179–191 (1999)
20. Steger, A., Wormald, N.C.: Generating random regular graphs quickly. Combinatorics, Probability & Computing 8(4), 377–396 (1999)